[ORBATOR] · CATEGORY INFLUENCE AUDIT #1

Influence Audit #1 — AI Tools for Lawyers

Who is steering the AI answers lawyers trust — measured.

ORBATOR · Category Influence Audit · Published 2026-08-10 · v1.2 Window: 28 days ending 2026-08-09 · 125 sampled answers · 5 engines (Claude, Gemini, ChatGPT, Perplexity, Grok) · 1,188 resolved citations Instrument disclosures: §7. Evidence: every number traces to retained answer texts + citation records.

Executive findings

1. **At most 34.9% of the sources AI consults in this category are independent** — and after human verification of the top-cited domains, at most 32.9%. Roughly **2 of every 3 citations resolve to domains owned by vendors being recommended in the answers themselves** (all vendor-owned shares in this report are floors; see §7).
2. **Every one of the top 22 most-cited pages is vendor-owned.** 20 belong to ranked vendors — and in every one of those cases, the page's owner was recommended in the same answer that cited it. The other 2 (eve.legal, gavel.io) belong to unranked category competitors whose "best legal AI" guides feature their own products, identified during pre-publication verification (§7). **The first independent source in the citation record is a Reddit thread, at rank 23.**
3. **A class of small vendors is recommended almost exclusively on their own word.** The Law GPT: self-cited in **13 of 13** runs where it was recommended (100%). 8am IQ for MyCase: **11 of 11** (100%). Definely: **7 of 7** (100%). AI Lawyer: 87.5%. LegesGPT: 85.7%. By contrast, category leaders earn recommendations mostly from other people's pages: Lexis+ 11.1%, Harvey 12.3%, Spellbook 21.3% self-sourced.
4. **Engines disagree materially.** Clio is recommended in 71% of Grok answers but 14% of ChatGPT's — a 57-point spread. Westlaw spans 37 points across engines, Spellbook and Harvey 36 each. "What AI recommends" depends heavily on *which* AI.
5. The pattern is structural, not episodic: measured independence was 28%, 29.6% and 34.9% across three successive windows (1,056 → 442 → 1,188 citations) — every window far more vendor-sourced than the AI-wide baseline.

---

1 · What AI actually recommends (measured, not editorial)

Share of the 125 sampled answers recommending each product (mention ≠ recommendation; Wilson 95% CI):

#	PRODUCT	REC. SHARE	95% CI
1	Spellbook (spellbook.com)	64.0%	55.3–71.9%
2	Harvey (harvey.ai)	58.4%	49.6–66.7%
3	CoCounsel (Thomson Reuters)	55.2%	46.5–63.6%
4	Microsoft Word (+Copilot)	53.6%	44.9–62.1%
5	Lexis+	50.4%	41.8–59.0%
6	Westlaw	45.6%	37.1–54.3%
7	Clio	33.6%	25.9–42.3%
8	vLex	28.8%	21.6–37.3%
9	LegalOn	27.2%	20.2–35.6%
10	Paxton / Vincent AI	25.6%	18.8–33.9%

Notable: the LexisNexis family (Lexis+, Lexis+ AI, LexisNexis) appears fragmented across product names; treated at family level it rivals the top spots. Microsoft Word's #4 position reflects engines recommending Word-with-Copilot for drafting — a reminder that AI answers include incumbents no legaltech listicle would rank.

2 · Citation independence: at most 34.9%

Of 1,188 resolved citations, at most 415 (34.9%) point to independent sources by registry matching against ranked vendors. The pre-publication verification pass reclassified two further top domains as vendor-owned (eve.legal, gavel.io — 24 citations between them), bringing measured independence to **at most 32.9%**. At least 67% of what AI reads in this category is the vendors' own web estate. For comparison, our 143k-citation global baseline is 46.2% vendor-owned — this category runs roughly **19 points more vendor-sourced than the AI-wide average**, placing it among the most vendor-steered categories we measure.

3 · The listicle economy (top cited pages)

The pages AI actually reads, ranked by citations; ownership resolved against the entity registry:

CITES	OWNERSHIP	PAGE
30	VENDOR — GC AI ★	gc.ai/blog/best-ai-tools-for-legal-research
28	VENDOR — Darrow ★	darrow.ai/resources/ai-tools-for-lawyers
21	VENDOR — Clio ★	clio.com/resources/ai-for-lawyers/ai-tools-for-lawyers
18	VENDOR — LegalOn ★	legalontech.com/post/best-ai-contract-review-tools
17	VENDOR — GC AI ★	gc.ai/blog/legal-ai-tools
17	VENDOR — Definely ★	definely.com/blogs/ai-contract-review-software
16	VENDOR — LEGALFLY ★	legalfly.com/post/best-ai-contract-review-software
16	VENDOR — Xantrion ★	xantrion.com/articles/best-ai-tools-for-lawyers-2026-comparison
15	VENDOR — Spellbook ★	spellbook.com/learn/legal-ai-tools
15	VENDOR — Eve (unranked competitor, verified)	eve.legal/blogs/best-legal-ai-software-lawyers
15	VENDOR — The Law GPT ★	thelawgpt.com/blog/best-ai-tools-contract-review-2026
13	VENDOR — AI Lawyer ★	ailawyer.pro/blog/10-ai-legal-research-tools
…	*(rows 13–22 all vendor-owned; full 25-row table in appendix)*	
9	INDEPENDENT — first in record, rank 23	reddit.com/r/legaltech — “best AI for contract analysis”

★ = the page’s owner was recommended in the same answer that cited the page. The dominant format is the vendor-authored “best AI tools for lawyers 2026” listicle. AI engines treat these as research; each one ranks its author. The two unranked-competitor pages (Eve’s and Gavel’s own “best/guide” listicles featuring their own products) were caught by human verification, not registry matching — evidence the true vendor share is higher than any automated count shows.

4 · Self-sourced recommendation rate (the dependency table)

Of the runs recommending each vendor, the share that also cite the vendor’s own domain:

VENDOR	SELF-SOURCED	READING
The Law GPT	**100%** (13/13 · small sample · CI 77–100%)	Every recommendation accompanied by its own pages
8am IQ for MyCase	**100%** (11/11 · small sample · CI 74–100%)	Same pattern
Definely	**100%** (7/7 · small sample · CI 65–100%)	Same pattern
AI Lawyer	**87.5%** (7/8 · small sample · CI 53–98%)	Near-total self-dependency
LegesGPT	**85.7%** (18/21 · CI 65–95%)	Near-total self-dependency
GC AI	**74.1%** (20/27)	
Darrow	66.7% (8/12)	
Clio	47.6% (20/42)	Rising — its own listicle is now the #3 cited page
LegalOn	44.1% (15/34)	
CoCounsel	34.8% (24/69)	Majority independently sourced
Spellbook	21.3% (17/80)	Recommendation earned mostly off-domain
Harvey	12.3% (9/73)	
Lexis+	11.1% (7/63)	

We report the measurement, not intent. But the split is stark: one class of vendors is recommended because the wider web talks about them; another is recommended, so far as the citation record shows, because they talk about themselves — and the engines listen.

5 · Engine divergence

Recommendation share by engine (top products; public Index view of the same window, n=111):

PRODUCT	CLAUDE	GEMINI	CHATGPT	PERPLEXITY	GROK†	SPREAD
Spellbook	52%	79%	50%	71%	86%	36pp
Harvey	71%	68%	50%	51%	86%	36pp
CoCounsel	62%	59%	43%	57%	71%	29pp
Westlaw	48%	50%	43%	34%	71%	37pp
Clio	19%	47%	14%	31%	71%	57pp

† Grok sample is small in-window (7 runs) — treat its column as indicative. Spreads of 29–57 points mean vendor visibility is engine-specific: a firm asking Gemini gets a different market than one asking ChatGPT. For buyers, this is why single-engine anecdotes mislead; for vendors, why engine-level measurement matters. Clio is the sharpest case: near-top visibility on Grok and Gemini, near-absent on ChatGPT.

6 · What this means (for each reader)

- **For lawyers/firms:** the "research" behind AI's tool advice is, two times out of three, the vendors' own marketing. Weight AI recommendations accordingly; ask for independent evidence.
- **For legaltech vendors:** the self-sourced table is the competitive map. A 100% self-sourced rival's position is fragile (it rests on sources engines may discount); a 12–21% self-sourced position is defensible. Both facts are actionable.
- **For the ecosystem:** this category has no heavily-cited independent comparison — the first independent source of any kind ranks 23rd. (That vacuum is what makes the listicle economy work.)

7 · Methodology & instrument disclosures

- 125 answers, 28d window ending 2026-08-09, 5 engines (Claude 28 runs, Gemini 34, ChatGPT 21, Perplexity 35, Grok 7); question set = neutral buyer intents; recommendation extracted under the mention≠recommendation rule (an entity merely named is not counted).
- Wilson 95% intervals on all shares; sample sizes shown; small-cell rows labeled.
- **Instrument notes, disclosed:** (1) All 5 engines measured current in-window (latest runs 2026-08-08/09); Grok's 7 runs are flagged small-sample. (2) Name-extraction was upgraded 2026-07-26; cross-window comparisons in this report are restricted to metrics insensitive to that change. (3) Registry hygiene performed before measurement: 6 duplicate entities merged (e.g. Harvey/Harvey AI, Paxton/Paxton AI), audit-logged. (4) §5's per-engine table is drawn from the public Index view of the same 4-week window (n=111; the metrics run includes one additional later batch). We disclose instrument changes in every report because sudden metric swings can otherwise be misread as market movement.
- Ownership resolution: citation domain → entity registry canonical domains (public record). **"Independent" = not owned by any vendor in this category — ranked or unranked.** Registry matching identifies ranked-vendor ownership automatically; unranked-competitor ownership is identified by human review of the top-cited domains. Because that review is finite, **all vendor-owned shares in this report are floors** — the true vendor-sourced share is at least the number stated. (The same floors caveat applies to our 143k-citation global baseline.)
- **Pre-publication verification (2026-08-05, refreshed 2026-08-10):** ownership of the top-cited domains was adversarially re-verified against the live web. Two domains counted independent by registry matching were reclassified as vendor-owned: eve.legal (Eve, plaintiff-law AI vendor; its cited page is its own "best legal AI" listicle) and gavel.io (Gavel, document-automation vendor; its cited "solo lawyers guide" features its own Gavel and Gavel Exec products). Both corrections strengthen finding #2.
- **Publish-time re-run:** all §1–§5 figures were recomputed on the current window at publication (2026-08-10) via the same instrument (influenceMetricsService); the prior draft's window (63 answers, 4 engines, 442 citations) produced materially consistent results (leader stable, independence within 5 points), disclosed here as the replication check.
- Evidence retention: full answer texts and citation records for every run in this report are retained and available for verification.

8 · Disclosure

Orbator operates the AI Recommendation Index and sells AI-visibility tooling. No vendor named in this report is an Orbator customer at publication. Orbator's own product categories are permanently excluded from this report series. Rankings and rates above are measured from sampled answers — never editorial, never paid. This report states measurements, not intent: "self-sourced" describes a citation pattern, not an allegation of wrongdoing.

Category Influence Audit #1 · Orbator Open Index · v1.2 · published 2026-08-10 · orbator.io/audits